AFRL-OSR-VA-TR-2014-0273

KLEINBERG COMPLEX NETWORKS

Robert Kleinberg
CORNELL UNIVERSITY

10/21/2014
Final Report

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTC
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
07-10-2014

2. REPORT TYPE
final

3. DATES COVERED *(From - To)*
Mar 2009 - Feb 2014

4. TITLE AND SUBTITLE
KLEINBERG COMPLEX NETWORKS

5a. CONTRACT NUMBER
FA9550-09-1-0100

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Kleinberg, Robert D.

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Cornell University

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Distribution A - Approved for Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT

This is the final report for AFOSR Award FA9550-09-1-0100. The award funded a research project conducted by PI Robert Kleinberg, investigating algorithmic aspects of the theory of complex networks as they relate to learning, routing, coding, and modeling of network structure and evolution. The research yielded a broad set of discoveries. Highlights over the five years of the project include the resolution of a twenty-year-old open question relating to the traveling salesman problem, the first rigorous analysis of Schelling's segregation process (a highly influential model that had been studied for decades via simulation in the social sciences, but whose behavior had resisted rigorous mathematical analysis prior to our work), a proof of the strongest separation to date between the power of linear and non-linear network codes, and a recent award-winning paper that introduces and analyzes a model of crowdsourced learning that sheds new light on the famous "exploration-exploitation" trade-off in sequential learning.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Robert D. Kleinberg
					19b. TELEPHONE NUMBER *(include area code)* 607-255-3313

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Final Report for AFOSR Award FA9550-09-1-0100

Robert Kleinberg

This report summarizes the research activities conducted under AFOSR Award FA9550-09-1-0100, which started on March 1, 2009 and concluded on February 28, 2014. The research addressed a range of problems relating to modeling, optimization, and learning in the context of complex networks. Within this broad area, the research can be organized into four distinct but inter-related threads: modeling of network structure and dynamics, network coding, optimization problems in network routing, and sequential learning in networks. This report is organized into four sections grouped around these themes. Within each section, we first present a summary of the project's achievements and then we report on the progress that was made during the most recent reporting period, which began on March 1, 2013.

Personnel. The grant supported three personnel during the course of the project.

1. The principal investigator, Robert Kleinberg, who led all aspects of the research.
2. Bruno Abrahao, a CS Ph.D. student at Cornell, was fully supported by this grant throughout the duration of the project. Abrahao received his Ph.D. in the summer of 2014 and is now a postdoctoral researcher at Stanford University, working under the supervision of Jure Leskovec.
3. Rad Niazadeh, a CS Ph.D. student at Cornell, was fully supported by this grant during the academic year 2012-13. Niazadeh participated in the research on sequential learning in networks before shifting to another project funded by a different grant.

1 Modeling of network structure and dynamics

The project made a number of contributions to our understanding of the structure of complex networks and the dynamics of stochastic processes, particularly epidemic processes, that take place on them. One highlight of this research, discussed in more detail below, was our analysis of Thomas Schelling's segregation process, a highly influential model that had been studied for decades via simulation in the social sciences, but whose behavior had resisted rigorous mathematical analysis prior to our work [13].

Much of the research on modeling network structure and dynamics was done by the PI in collaboration with Bruno Abrahao, a Ph.D. student who was fully supported by the grant. Shortly before the start of this project, we had co-authored a paper on the fractal nature of the Internet's delay space, the matrix of measured packet round-trip times between hosts on the Internet [4].

We had originally envisioned that a significant portion of the project would involve broadening this research to incorporate new datasets describing distance estimates in other complex networks, while developing an accompanying body of theoretical work relating to structural characteristics of complex metric spaces and algorithmic problems relating to the task of embedding them in simpler "host metrics" such as low-dimensional normed vector spaces. However, as the project evolved we shifted our emphasis to a set of modeling and inference questions inspired by the analysis of social networks rather than communication networks. A few different factors contributed to this shift in emphasis. First, we anticipated that the most important future research challenges in Internet measurement and modeling would come not from the low-level communications infrastructure but from the high-level processes (e.g. social networking) that are built upon this infrastructure. In hindsight, I believe this was an accurate prediction of how the field evolved over the past five years. Second, it was difficult to obtain new large-scale datasets of Internet distance measurements, due to changes beyond our control in the way servers are configured worldwide. In contrast, publically available social network datasets become richer and more plentiful with each passing year.

The PI's research with Abrahao on social network structure and dynamics confronted two questions. First, to what extent is the presence of "communities" in a network reflected in the network's link structure and how can we quantify the effectiveness of community detection algorithms? Second, how much data do we need about the spread of epidemic cascades on a network in order to reconstruct the network's adjacency matrix? Our research on the first question [2, 3] led to the development of a novel "reverse classification" paradigm for using a supervised classifier (a type of machine learning algorithm) in order to detect and quantify similarities between entire classes of labeled examples, as opposed to the traditional use of such classifiers for assigning new labels to unlabeled data. Our results called into question established assumptions about the structure of communities in online social networks, suggesting that naïve baseline algorithms (variants of random walks) extracted subsets whose structure more closely resembled annotated communities than the structure of subsets extracted by sophisticated algorithms commonly used by network community detection researchers. Our research on the second question [1] yielded rigorous upper and lower bounds on the amount of data required for network inference, i.e., the task of reconstructing the edge set of a network from traces of independent epidemic cascades randomly propagating on that network. We also provided a new algorithm whose precision and recall outperformed that of the best previous network inference algorithm, for the case in which a moderate or large number of traces are available.

In addition to the PI's collaboration with Bruno Abrahao, the project involved several other collaborations that yielded new findings on random processes in networks. With economists Larry Blume and David Easley, and computer scientists Jon Kleinberg and Éva Tardos, we investigated so-called *threshold contagion processes* in networks and their influence on the formation of the network itself. In a series of papers [10, 11, 12] we study both the pure optimization aspect of this problem — how to design a network to minimize the spread of contagion — and its strategic aspects: which networks will be formed by selfish individuals seeking to balance the benefit of linking to peers against their own risk of infection?

In joint work with Jon Kleinberg, Seth Marvel, and Steven Strogatz [22] we studied a different type of process that plays out over time in networks: the shifting patterns of alliance and conflict among a set of interacting parties. Sociologists have long observed that such patterns tend to satisfy a property called *structural balance*. Our work provides a mathematical model in which this observation can be rigorously justified: we investigated a simple differential equation in which the strength of the tie between two individuals is represented by a real number (positive for friendship, negative for conflict) and these numbers increase or decrease over time according to the strength of common friends or enemies between the two individuals. We proved that for most initial conditions, the matrix of tie strengths (suitably rescaled) converges to a rank-one matrix; this explains the structural balance property, since sign patterns of rank-one matrices are easily seen to satisfy structural balance. Our work appeared in Proceedings of the National Academy of Sciences (PNAS) and attracted some press attention, including an article in the *Los Angeles Times* [14].

Finally, in joint work with Christina Brandt, Gautam Kamath, and Nicole Immorlica [13], we investigated Schelling's segregation model, a famous model of residential segregation that has been observed, in simulations, to yield substantially segregated patterns even when no individual in the simulation has a strict preference for living in a segregated neighborhood. Our work provides the first mathematically rigorous justification for this observation, while showing that the extent of segregation implied by Schelling's model, as one scales the number of individuals and the population density, may be less dramatic than is commonly believed. Once again, this work was picked up by the press [18].

1.1 Work performed in the most recent reporting period

During the most recent reporting period, we finalized the journal versions of two papers cited above [3, 11]. The first of these is the paper on characterizing community structure in networks, the second is the paper on models of network formation by self-interested agents in the presence of contagious risk.

Final versions of two more papers [1, 13] are currently in preparation. In the course of revising these papers for journal publication, we have discovered new results that significantly strengthen the original work. First, we developed a new information-theoretic technique for proving lower bounds on the trace complexity of network inference that is significantly stronger and more versatile than the technique presented in [1], producing sharper quantitative results and even yielding enhanced lower bounds that apply to the task of approximately (rather than exactly) reconstructing the network's edge set. Second, we strengthened our technique for analyzing Schelling's segregation process to show that the expected average length of monochromatic runs in the final labeling is $O(w)$, improving the $O(w^2)$ bound reported in [13]. This improvement is significant since it matches the $\Omega(w)$ lower bound presented in [13], meaning that our analysis of one-dimensional Schelling segregation is now tight up to constant factors.

2 Network coding

Given a network consisting of links with limited bandwidth, how can we tell if a given set of communication streams can be transmitted simultaneously without violating the network's bandwidth limitations? This fundamental information theoretic question, which is central to the theory of network coding, has proven surprisingly resistant to the efforts of scores of researchers worldwide. As one indication of how little is known about the problem in general, we still do not know of any algorithm — even one with exponential running time — to compute a non-trivial (asymptotically better than $O(n)$) approximation to the maximum concurrent network coding rate in multiple unicast networks. One highlight of the network coding research conducted under this project was the design and analysis of the first algorithm with a non-trivial approximation ratio for the important special case of broadcasting with side information problems, which correspond to the special case of multiple unicast networks in which only one edge has finite capacity.

The PI's research on this portion of the project was joint work with Cornell Ph.D. student Anna Blasiak, who was funded by a separate fellowship, and Eyal Lubetzky, a mathematician at Microsoft Research. Our work focused on communication problems in which streams of information are transmitted using a broadcast channel to receivers, each of whom wants to decode one or more of the streams and potentially has side information revealing the data in some of the other streams. Given the side information mapping, the goal is to determine the minimum broadcast channel capacity that allows successful transmission of the desired streams to all parties. Prior work had explored relationships between the broadcast rate and graph-theoretic parameters such as the independence number and the clique-cover number of the side information graph. Our work [9] successfully applied an information-theoretic linear program to derive tight bounds on the broadcast rate for several families of graphs, the first examples of graphs whose broadcast rate could be precisely determined except for those cases in which it coincided with the independence number and the clique-cover number. In another paper [8] we introduced a lexicographic product operation that combines two index coding instances into a larger one, and we showed that the broadcast rate is sub-multiplicative under this operation while our linear-programming-based lower bound is super-multiplicative. This technical result enabled us to apply the new product operation to amplify the magnitude of certain gaps between combinatorial and information-theoretic parameters of networks. For example, we showed that the best known combinatorial upper bound on the network coding rate for multiple unicast networks — the *informational meagerness* [17] — could exceed the actual network coding rate by a factor that grows polynomially in the network size, an exponential improvement over the best previously known separation.

Probably the most significant contribution of our work was a strong separation between vector-linear and non-linear network coding. As the most versatile known general-purpose network coding method, vector-linear codes have been used intensively in the network coding literature. Prior to our work, it had been shown by Dougherty, Freiling, and Zeger that non-linear coding can outperform vector-linear coding by a factor of at least 1.1 in some networks, but no stronger separation was known, leading some researchers to believe that vector-linear codes may achieve near-optimal rates for network coding problems in general. Our work [8] showed this is not possible, by applying our lexicographic product construction to furnish examples of networks in which the broadcast rate achievable by non-linear network codes is superior to the vector-linear broadcast rate by a factor of n^ϵ for an absolute constant $\epsilon > 0$.

2.1 Work performed in the most recent reporting period

Our first paper on index coding was accepted for publication in *IEEE Transactions on Information Theory* during the most recent reporting period [9].

The PI was no longer actively conducting new research on network coding during the most recent reporting period, but he was still supervising the research of Anna Blasiak. In 2013 she used network coding theory — specifically, an analysis of the behavior of network coding rates under a novel type of graph product operation, different from the lexicographic product that we analyzed in [8] — to re-derive a lower bound on the multiplicative separation between the maximum multicommodity flow and minimum multicut values in directed graphs, originally due to Saks and Samorodnitsky. Blasiak's method achieves a slight quantitative improvement over the Saks-Samorodnitsky result, yielding an exact determination of the minimum multicut value for their network. Her work [7] received the student paper award at NetCod 2013.

3 Optimization problems in network routing

Many of the most fundamental optimization problems in network theory are routing problems, i.e. problems that require optimizing over paths or linear combinations of paths. This project featured research on two classes of routing problems, namely traveling salesman problems and multicommodity flows. One highlight of this research was our discovery of a polynomial-time algorithm for the metric traveling salesman s-t path problem whose approximation ratio achieves the first provable improvement upon the approximation ratio of the famous algorithm introduced by Christofides in 1976 (which was first analyzed in the context of the s-t path variant by Hoogeveen in 1991). A corresponding improvement for the more familiar cycle variant of the metric TSP would resolve one of the most venerable open problems in the theory of approximation algorithms.

Our research on traveling salesman problems made use of a familiar linear programming relaxation of the problem known as the Held-Karp or subtour relaxation. The optimal integer solution of the Held-Karp linear program is a minimum-cost traveling salesman tour (or s-t path, if the appropriate variant of the linear program is used). In [5] we analyzed a simple polynomial-time algorithm for rounding a fractional solution of the Held-Karp linear program to an integer solution; the algorithm simply represents the fractional solution as a linear combination of (at most polynomially many) spanning trees. It transforms each tree into a spanning subgraph that contains an Eulerian s-t path, by using a minimum-cost matching to correct the vertices whose degree has the wrong parity. Then it "shortcuts" the Eulerian path in each of these spanning subgraphs to obtain a Hamiltonian s-t path and takes the cheapest path thus obtained. We proved that this rounding procedure never increases the cost of the fractional solution by a factor greater than the golden ratio, $1.618\ldots$, thus improving the bound of 5/3 due to Hoogeveen.

Earlier in the project, we had analyzed the same linear programming relaxation of the metric *asymmetric* traveling salesman problem, proving that the Held-Karp relaxation yields a $O\left(\frac{\log n}{\log\log n}\right)$ approximation for the objective of minimizing the maximum edge cost, the so-called *bottleneck* objective. Our paper [6] leveraged a maximum-entropy rounding technique introduced in an earlier paper by Asadpour et al. that had obtained the same approximation ratio with respect to the sum objective, for which the best previously known bound was $O(\log n)$. Surprisingly, for the bottleneck objective no algorithm was previously known to achieve an approximation ratio asymptotically better than the trivial $O(n)$ approximation. Another variant of the metric TSP that we analyzed earlier in the project is the so-called universal TSP problem, in which one must compute a total ordering of the points (a so-called *universal tour*) whose induced tour on each subset of the metric space approximates the minimum-cost TSP tour of that subset. Our work [16] improved the best known lower bound for this problem to $\Omega(\log n)$, which is most likely tight (up to constant factors) although the best upper bound known at present is $O(\log^2 n)$.

Optimization problems in network routing are at the core of the Merlin project [23], a collaboration with colleagues at Cornell University and the University of Lugano, which combines research on programming languages, distributed systems, and algorithms into a system that provides an improved interface and algorithms for network management and traffic engineering. At a technical level, the PI's algorithmic contribution to Merlin was a construction that allows multi-commodity flow algorithms to perform optimal load-balancing subject to edge capacity constraints and user-specified *path constraints* stipulating that packets may only traverse paths that satisfy a designated regular expression over the alphabet consisting of the addresses of all devices in the network. Traditional multi-commodity flow algorithms tolerate edge capacity constraints but not path constraints; the extension to path constraints is necessary in order for network operators to express policies consistent with their security, functionality, and traffic engineering goals.

3.1 Work performed during the most recent reporting period

An expanded version of our paper on the metric traveling salesman s-t path problem [5] was accepted for publication in *Journal of the ACM*, pending revision, during the most recent reporting period.

During this reporting period we commenced participation in the Merlin project, designing and applying algorithms for the variants of multicommodity flow that arise in programmable network management. This research has been presented at the *HotNets 2013* and *CoNEXT 2014* conferences and several other networking workshops and meetings, and it has attracted the interest of many groups in industry contemplating the use of similar systems for traffic engineering in their wide-area or datacenter networking operations.

4 Sequential learning in networks

Sequential learning theory studies the design of optimal prediction and decision rules, in dynamic settings where the learner's action at time t may be based on the data observed at all earlier times. Such problems frequently incorporate some version of the "exploration-exploitation dilemma", in which the myopically optimal action may not provide the most informative data for the future. While optimal learning policies are well understood in the simplest cases (such as the famous multi-armed bandit problem), in many cases of interest the optimal policy suffers from prohibitive complexity and it is necessary to turn to designing approximately optimal policies. Our research on this subject focused on extensions of the multi-armed bandit problem in which there are side constraints relating the "arms" to one another. Such side constraints arise commonly in networked settings, where the network structure enforces correlations among observations and/or payoffs.

One highlight of our research on these problems was a paper [15] that studied "crowdsourced learning" via a game-theoretic variant of multi-armed bandits in which arms are pulled by selfish and myopic agents who may be induced to play actions other than the myopically optimal one using suitable side payments. Our work, which introduced this model and precisely quantified the exploration-exploitation tradeoff in this setting, received the best paper award at the 2014 ACM Conference on Economics and Computation.

Earlier in the project, we collaborated with Aleksandrs Slivkins and Eli Upfal on multi-armed bandit problems in which there is a large (potentially infinite) set of arms, but with side constraints that place upper bounds on the absolute difference in payoff between certain pairs of arms. This induces a metric space structure on the arms. Our work [20] showed that the performance of the optimal policy, as measured by its *regret* (expected payoff relative to the *a posteriori* optimal arm) is determined by geometric parameters related to covering numbers of this metric space. Intriguingly, there is a dichotomy between metric spaces admitting algorithms with near-logarithmic regret, and those in which the regret of the optimal policy grows at a rate of $O(t^{1/2})$ or worse; no exponent between 0 and 1/2 is possible. Our work [19] exposes this learning-theoretic dichotomy and reveals it to be a manifestation of a much older, purely topological dichotomy due to Cantor and Bendixson.

In another recent paper [21], joint with Bruno Abrahao and with co-authors at Microsoft Research Asia, we studied so-called *combinatorial partial monitoring problems*, a class of learning problems derived from networked settings in which the learner's actions result in partial observations of an underlying network state that is not fully observed. For example, choosing a routing path in a network may reveal the delay on all of the edges in the path but not the delays on other, unobserved edges of the network. The potential correlations among the observable variables associated to the different actions exacerbate the complexity of the problem. We were able to design a learning algorithm whose regret grows at a rate which is sublinear in general, and polylogarithmic assuming a sufficiently large gap between the optimal and second-best alternatives.

4.1 Work performed during the most recent reporting period

Our work on crowdsourced learning [15] was done during the most recent reporting period, and a journal version is now in preparation. Our work on combinatorial partial monitoring [21] was also done during the most recent reporting period.

References

[1] B. Abrahao, F. Chierichetti, R. Kleinberg, and A. Panconesi. Trace complexity of network inference. In *Proceedings of the 19th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 491–499, 2013.

[2] B. Abrahao, S. Soundarajan, J. Hopcroft, and R. Kleinberg. On the separability of structural classes of communities. In *Proceedings of the 18th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 624–632, 2012.

[3] B. Abrahao, S. Soundarajan, J. Hopcroft, and R. Kleinberg. A separability framework for analyzing community structure. *ACM Transactions on Knowledge Discovery from Data*, 8(1), 2013. Special issue on Computational Aspects of Social and Information Networks.

[4] B. D. Abrahao and R. D. Kleinberg. On the Internet delay space dimensionality. In *Proceedings of the 8th ACM SIGCOMM Conference on Internet Measurement (IMC)*, pages 157–168, 2008. Brief announcement appeared in *Proceedings of the Twenty-Seventh Annual ACM Symposium on Principles of Distributed Computing, PODC 2008* .

[5] H.-C. An, R. Kleinberg, and D. B. Shmoys. Improving Christofides' algorithm for the s-t path TSP. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*, pages 875–886, 2012.

[6] H.-C. An, R. D. Kleinberg, and D. B. Shmoys. Approximation algorithms for the bottleneck asymmetric traveling salesman problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 13th International Workshop, APPROX 2010, and 14th International Workshop, RANDOM 2010*, volume 6302 of *Lecture Notes in Computer Science*, pages 1–11. Springer, 2010.

[7] A. Blasiak. Multicut lower bounds via network coding. In *Proceedings of the 2013 IEEE Symposium on Network Coding (NetCod)*, 2013.

[8] A. Blasiak, R. Kleinberg, and E. Lubetzky. Lexicographic products and the power of non-linear network coding. In *Proceedings of the 52nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 609–618, 2011.

[9] A. Blasiak, R. Kleinberg, and E. Lubetzky. Broadcasting with side information: Bounding and approximating the broadcast rate. *IEEE Transactions on Information Theory*, 59(9):5811–5823, 2013.

[10] L. Blume, D. A. Easley, J. M. Kleinberg, R. D. Kleinberg, and É. Tardos. Network formation in the presence of contagious risk. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pages 1–10, 2011.

[11] L. Blume, D. A. Easley, J. M. Kleinberg, R. D. Kleinberg, and É. Tardos. Network formation in the presence of contagious risk. *ACM Transactions on Economics and Computation*, 1(2), 2013.

[12] L. E. Blume, D. A. Easley, J. M. Kleinberg, R. Kleinberg, and É. Tardos. Which networks are least susceptible to cascading failures? In *Proceedings of the 52nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 393–402, 2011.

[13] C. Brandt, N. Immorlica, G. Kamath, and R. Kleinberg. An analysis of one-dimensional Schelling segregation. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*, pages 789–804, 2012.

[14] E. Brown. An equation for friendship. *Los Angeles Times*, January 14 2011. http://www.latimes.com/health/la-sci-friend-or-foe-20110115,0,763300.story.

[15] P. Frazier, D. Kempe, J. Kleinberg, and R. Kleinberg. Incentivizing exploration. In *Proceedings of the 15th ACM Conference on Electronic Commerce (EC)*, 2014.

[16] I. Gorodezky, R. D. Kleinberg, D. B. Shmoys, and G. Spencer. Improved lower bounds for the universal and *a priori* TSP. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 13th International Workshop, APPROX 2010, and 14th International Workshop, RANDOM 2010, Barcelona, Spain, September 1-3, 2010. Proceedings*, volume 6302 of *Lecture Notes in Computer Science*, pages 178–191. Springer, 2010.

[17] N. J. A. Harvey and R. Kleinberg. Tighter cut-based bounds for k-pairs communication problems. In *Proceedings of the 43rd Annual Allerton Conference on Communication, Control, and Computing*, 2005.

[18] B. Hayes. Computing Science: The Math of Segregation. *American Scientist*, 101(5): 338–341, 2013.

[19] R. Kleinberg and A. Slivkins. Sharp dichotomies for regret minimization in metric spaces. In *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 827–846, 2010.

[20] R. Kleinberg, A. Slivkins, and E. Upfal. Multi-armed bandits in metric spaces. In *Proceedings of the 39th ACM Symposium on Theory of Computing (STOC)*, pages 681–690, 2008.

[21] T. Lin, B. Abrahao, R. Kleinberg, J. C. S. Lui, and W. Chen. Combinatorial partial monitoring game with linear feedback and its applications. In *Proceedings of the 31st International Conference on Machline Learning (ICML)*, 2014.

[22] S. A. Marvel, J. Kleinberg, R. D. Kleinberg, and S. H. Strogatz. Continuous-time model of structural balance. *Proceedings of the National Academy of Sciences (PNAS)*, 108(5):1771–1776, 2011.

[23] R. Soulé, S. Basu, P. Jalili Marandi, F. Pedone, R. Kleinberg, E. G. Sirer, and N. Foster. Merlin: A language for provisioning network resources. In *Proceedings of the 10th Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, 2014. Preliminary version in *Proceedings of the 12th ACM Workshop on Hot Topics in Networks (HotNets-XII)*, 2013.